

LISTING OF THE CLAIMS

1. (Previously Presented) An apparatus for a computer system, comprising:
a prefetcher coupled to a first memory for a processor of the computer system, the first memory having a first latency; and
a tracker within the prefetcher and configured to observe accesses by a processor to a plurality of cache lines wherein the accesses form a stream-type sequential access pattern having a direction that is tracked by setting bits in a bit vector, and wherein further the tracker is configured to use the bit vector to predictively load a target cache line indicated by the direction of the stream-type sequential access pattern from the first memory into a second memory having a second latency less than the first latency.

2. (Previously Presented) The apparatus of claim 1, wherein the tracker includes a tag configured to recognize accesses to corresponding cache lines of the first memory by the processor.

3. (Previously Presented) The apparatus of claim 2, wherein a plurality of accesses by the processor to the first memory as recognized by the tag are used by the tracker to determine the target cache line for a predictive load into the second memory.

4. (Previously Presented) The apparatus of claim 3, wherein consecutive accesses by the processor to adjacent cache lines of the first memory are used to determine the target cache line for a predictive load into the second memory, and wherein the adjacent cache lines have adjacent addresses.

5. (Previously Presented) The apparatus of claim 1, wherein the first memory comprises a memory block of a plurality of memory blocks of the computer system.

6. (Previously Presented) The apparatus of claim 5, wherein the first memory comprises a four kilobyte page of system memory of the computer system.

7. (Previously Presented) The apparatus of claim 5, wherein the tracker includes a tag configured to monitor a sub portion of the first memory for accesses by the processor.

8. (Previously Presented) The apparatus of claim 1, wherein the first memory is a system memory of the computer system.

9. (Previously Presented) A system, comprising:
a processor;
a first memory coupled to the processor, wherein the first memory has a first latency;
a prefetch unit coupled to the first memory;
a plurality of trackers included in the prefetch unit, wherein the trackers are respectively configured to recognize processor accesses to the first memory and to observe accesses to cache lines within a second memory having a second latency less than the first latency, the second memory operable to supply data to the processor responsive to processor data requests that form a sequential access pattern comprising an order in which adjacent storage locations in the first memory are accessed, wherein the order is tracked by setting bits in a bit vector; and
the second memory coupled to the prefetch unit, wherein the prefetch unit uses the bit vector to predictively load target cache lines from the first memory into the cache

memory, and wherein the target cache lines are indicated by the sequential access pattern identified by the bit vector.

10. (Previously Presented) The system of claim 9, wherein each of the trackers include a tag to recognize accesses to cache lines by the processor.

11. (Previously Presented) The system of claim 9, wherein a plurality of first memory accesses by the processor are used by the trackers to determine the target cache lines for a predictive load into the second memory.

12. (Previously Presented) The system of claim 11, wherein consecutive accesses by the processor to adjacent cache lines of a page are used to determine the target cache line for a predictive load into the second memory, wherein the adjacent cache lines have adjacent addresses.

13. (Previously Presented) The system of claim 9, wherein the first memory comprises a plurality of 4 KB pages.

14. (Previously Presented) The system of claim 9, wherein each of the plurality of trackers is configured to monitor a sub-portion of a page for accesses by the processor.

15. (Previously Presented) The system of claim 14, wherein the cache lines are 128 byte cache lines and wherein a tag is used to monitor half of a page for accesses by the processor.

16. (Previously Presented) The system of claim 9, wherein the second memory is a prefetch cache memory within the prefetch unit.

17. (Previously Presented) The system of claim 9, wherein the second memory is an L2 cache memory.

18. (Previously Presented) A method, comprising:
monitoring data transfers between a first memory having a first latency and a second memory coupled to a processor by using a prefetcher, wherein the first memory has a first latency and the second memory has a second latency less than the first latency, and wherein the prefetcher is coupled to the first memory, wherein the second memory is a cache memory operable to supply data to the processor responsive to processor data requests;

using bit vectors to track multiple stream-type sequential access patterns by the processor to the first memory, wherein the access patterns comprise orders in which adjacent storage locations in the first memory are accessed, wherein the orders are tracked by setting bits in respective bit vectors; and

prefetching data from the first memory to the second memory as indicated by the access patterns identified by the bit vectors.

19. (Previously Presented) The method of claim 18 wherein the computer system comprises a plurality of processors, and wherein the processors include a processor that is coupled to a first memory and a second memory.

20. (Previously Presented) The method of claim 18, wherein consecutive accesses by the processor to adjacent cache lines of the first memory are used to

determine a target cache line of a stream type access pattern for a prefetching to the second memory, wherein the adjacent cache lines have adjacent addresses, and wherein the target cache line is part of stream-type accesses that formed the stream-type access pattern.

21. (Previously Presented) The system of claim 9, wherein said prefetch unit accesses to first memory are timed to utilize processor-to-system memory idle time.

22. (Previously Presented) A device, comprising:
means for observing data transfers between a first memory, having a first latency, and a second memory coupled to a processor, wherein the second memory is a cache memory operable to supply data to the processor responsive to processor data requests, and wherein the second memory has a second latency less than the first latency;
means for using bit vectors to track multiple stream-type sequential access patterns by the processor to the first memory, wherein the access patterns comprise orders in which adjacent storage locations in the first memory are accessed, wherein the orders are tracked by setting bits in respective bit vectors; and
means for prefetching data from the first memory to the second memory as indicated by the access patterns identified by the bit vectors.

23. (Previously Presented) The device of claim 22 wherein the computer system includes a plurality of processors, and wherein each of the processors is coupled to a respective first memory and a second memory.

24. (Previously Presented) The device of claim 22, wherein consecutive accesses by the processor to adjacent cache lines of the first memory are used to determine a target

cache line of a stream type access pattern for a prefetching to the second, wherein the adjacent cache lines have adjacent addresses, and wherein the target cache line is part of stream-type accesses that formed the stream-type access pattern.

25. (Previously Presented) The apparatus of claim 1, wherein the prefetcher comprises a prefetch cache operable to be used to load a cache line from the first memory.